

On product, generic and random generic quantum satisfiability

C. R. Laumann,¹ A. M. Läuchli,² R. Moessner,² A. Scardicchio,³ and S. L. Sondhi⁴

¹*Department of Physics, Princeton University, Princeton, NJ 08544, USA*

²*Max Planck Institut für Physik Komplexer Systeme, 01187 Dresden, Germany*

³*Abdus Salam International Centre for Theoretical Physics, Strada Costiera 11, 34014 Trieste, Italy*

⁴*Department of Physics, Princeton University, Princeton, NJ 08544*

(Dated: 1 July 2010)

We report a cluster of results on k -QSAT, the problem of quantum satisfiability for k -qubit projectors which generalizes classical satisfiability with k -bit clauses to the quantum setting. First we define the NP-complete problem of product satisfiability and give a geometrical criterion for deciding when a QSAT interaction graph is product satisfiable with positive probability. We show that the same criterion suffices to establish quantum satisfiability for *all* projectors. Second, we apply these results to the random graph ensemble with generic projectors and obtain improved lower bounds on the location of the SAT–unSAT transition. Third, we present numerical results on random, generic satisfiability which provide estimates for the location of the transition for $k = 3$ and $k = 4$ and mild evidence for the existence of a phase which is satisfiable by entangled states alone.

Contents

I. Introduction	1
II. Product satisfiability on general graphs	2
A. Counting product states at $M = N$	4
III. Application to random graphs	4
IV. Beyond PRODSAT	5
A. Direct search for SAT–unPRODSAT instances	5
B. Finite size scaling	6
V. In closing	6
Acknowledgements	7
A. Formal PRODSAT proof	7
B. Reduced density matrix diagnostic for product states	7
C. Improved estimate of α_c^+ from BMR sunflowers	8
References	9

I. INTRODUCTION

The central achievement of classical complexity theory was the discovery that certain natural classes of constraint satisfaction problems, prototypically boolean satisfiability (SAT), can encode all verifiable decision problems in an efficient manner. This insight, encapsulated in the Cook-Levin theorem, and the general consensus that $P \neq NP$, imply that any classical approach to solving instances of SAT must in general take exponentially long [1]. Unfortunately, the recent development of an analogous set of quantum complexity theoretic results

suggest that a similar situation holds even for quantum computers: in this case, quantum satisfiability (QSAT) can encode any problem in the class QMA₁[15] and ought therefore resist solution by even the most subtle quantum attacks [2].

These hardness results, while rigorous, are in some sense unsatisfactory. They imply that (quantum) satisfiability is exponentially hard to solve in the worst-case, but shed little light on what separates the worst-case instance from the trivial. In this pursuit, it is useful to introduce a measure on instances of satisfiability and study the statistical features of the ensemble rather than the worst-case behavior of the whole class. This approach provides control parameters with which to tune our attention across various regimes of difficulty and also allows us to bring to bear the full arsenal of statistical mechanics and probability theory [3, 4].

This paper is devoted to such a study in the context of *quantum* complexity theory. We study quantum satisfiability (k -QSAT), the decision problem of determining whether a given Hamiltonian H acting on N qubits has a ground state with precisely zero energy. The corresponding Hamiltonian has the form

$$H = \sum_{m=1}^M \Pi^m \quad (1)$$

where Π^m is a k -local (*i.e.* k -qubit) projector associated with the (hyper-)edge $m = (m_1, m_2, \dots, m_k)$ of an interaction graph G . Such a Hamiltonian has a ground state $|\psi\rangle$ of zero energy if and only if $|\psi\rangle$ is simultaneously annihilated by all of the Π^m . Furthermore, an instance of k -QSAT must come with a promise $\delta = 1/\text{poly}(N)$ that the ground state energy of H , should it be greater than zero, will actually be greater than δ . We will have little more to say regarding the promise gap in this paper, although we believe that it is satisfied in the thermodynamic limit of the ensembles we consider.

In Bravyi's original work introducing the k -QSAT

problem, it was shown to be QMA_1 -complete[2] for $k \geq 4$. In a previous paper[5], we introduced a random ensemble for k -QSAT and worked out the coarse features of its phase diagram in the large N limit. There are two sources of randomness in the k -QSAT ensemble: 1. the discrete choice of interaction graph G and 2. the continuous choice of projectors Π^m associated to each edge. The latter choice is particularly powerful: *generic* choices of projectors reduce quantum satisfiability to a graph, rather than Hamiltonian, property. More precisely, for fixed G , the dimension $R_G = |\ker H|$ of the satisfying subspace of H is *almost always* minimal with respect to the choice of projectors Π^m . Thus, if G can be frustrated by *some* choice of projectors, it is frustrated for almost all such choices. This “geometrization” property allows us to make strong statements about the quantum satisfiability of Hamiltonians associated with both random and non-random graphs and even non-generic choices of projector, both analytically and numerically. Throughout this paper, we use the term *generic* to refer to the continuous choice of projectors and *random* to refer to the choice of the graph.

This paper presents a cluster of results regarding *generic* quantum satisfiability under the restriction that the projectors $\Pi^m = |\phi\rangle\langle\phi|$ are rank 1. In Sec. II we introduce and characterize product satisfiability (PRODSAT) on non-random graphs G . This essentially classical property is NP-complete, but will allow us to prove that interaction graphs in which each clause (edge) may be matched with a neighboring qubit (node) are quantum satisfiable for *any* choice of projectors. For the marginal case $M = N$, in which these correspond to perfect matchings of the clause-qubit graph, the associated product states are discrete and provide some access to the ground state space, as will be explored in Sec. II A. As physicists, we refer to the qubit-clause matchings as *dimer coverings*, although computer scientists will likely prefer to think of perfect matchings and Hall’s marriage condition.

The remainder of the paper is especially concerned with mapping out the phase diagram of generic quantum satisfiability on large *random* interaction graphs. As in our previous work[5], we focus on the Erdős-Renyi ensemble with clause density α , in which each of the potential $\binom{N}{k}$ edges appears with probability $p = \alpha N / \binom{N}{k}$. As one might expect, the ensemble enjoys a satisfiable (SAT) phase at low clause densities α and becomes frustrated (unSAT) at sufficiently high clause density. Originally, we showed the existence of the SAT phase by constructing satisfying product states using something akin to a transfer matrix approach. This algorithm works generically so long as the graph G does not have a hypercore – that is, a maximal subgraph in which every qubit participates in at least two clauses. This property is guaranteed for $\alpha < \alpha_{hc}(k)$, the critical density beyond which a hypercore appears. These critical values are available in the literature[6]; for $k = 3$, $\alpha_{hc} \approx 0.81$, while for $k \rightarrow \infty$, $\alpha_{hc} \rightarrow 0$.

In Sec. III, we apply the characterization of product

state satisfiability of Sec. II to the random graph ensemble. This improves the lower bound on the SAT/unSAT transition from α_{hc} to α_{dc} , the critical value for the existence of dimer coverings. For $k = 3$, $\alpha_{dc} \approx 0.92$ is not much of an improvement over $\alpha_{hc} \approx 0.81$, but as $k \rightarrow \infty$, $\alpha_{dc} \rightarrow 1$ instead of 0. Moreover, this is a complete characterization of the phase in which the ensemble is generically satisfiable by product states. If there is a satisfiable phase for larger α , its satisfying states will generically be entangled.

At high clause densities, a simple local bound on the generic dimension of the zero energy subspace provided the first evidence for the existence of an unSAT phase. More recently, Bravyi *et al* (BMR) [7] have dramatically improved this bound by finding the ground state degeneracy of larger clusters (“sunflowers” and “nosegay” graphs) and determining their prevalence within the large random graph. Using these techniques they have placed an upper bound, α_c^+ , on the SAT/unSAT transition of $\alpha_c^+ \approx 3.594$ for $k = 3$ and with an exponential scaling $\alpha_c^+ \sim 2^k$ for large k .

The gap between the best rigorous lower and upper bounds on the SAT/unSAT transition is distressingly large: $\alpha_c^- = \alpha_{dc} = O(1)$ versus $\alpha_c^+ = O(2^k)$. Moreover, the existence of a satisfiable regime above α_{dc} would provide a large ensemble of large graphs with intrinsically entangled satisfying states. Thus, in Sec. IV, we address the situation for larger values of α by adducing some primarily numerical evidence about the existence of random instances which are SAT but not PRODSAT. We note that recently, Ambainis *et al* [8] have proven the existence of such an entangled SAT phase for $k \geq 12$ using their newly proven quantum Lovász local lemma coupled with the characterizations presented in this paper. We will have more to say about this exciting development in the conclusion.

II. PRODUCT SATISFIABILITY ON GENERAL GRAPHS

One of the remarkable results in Bravyi’s original work on QSAT was the explicit construction of satisfying product states for $k = 2$, thanks to which 2-QSAT turned out to be easy to decide and the concomitant satisfying states easy to construct. This poses the question whether there exist regimes for $k \geq 3$ -QSAT in which analogous results hold. We therefore define:

Definition 1. *An instance of QSAT, (G, Π^m) , is PRODSAT if it has a satisfying product state. That is, $\exists |\Psi\rangle = |\psi_1\rangle \otimes \cdots \otimes |\psi_N\rangle$ s.t. $\Pi^m |\Psi\rangle = 0 \ \forall m \in G$.*

PRODSAT may be viewed as a decision problem in its own right and, in the presence of a polynomially small promise gap δ , it is efficiently verifiable by classical computation. A witness is simply a collection of $2N$ \mathbb{C} components of a satisfying product state, the energy of which




FIG. 1: (Color online) (a) Example of a $k = 3$ interaction graph with $M < N$. Circles (green) indicate qubits and squares (red) indicate clause projectors that act on adjacent qubits; (b) a dimer (shaded blue) covering that covers all clauses.

may be evaluated in linear time. Moreover, by choosing Π_m to project onto computational basis states, it is clear that PRODSAT contains SAT and is therefore NP-complete.

In this section, we investigate product satisfiability for fixed graphs G and general choices of rank 1 projectors $\Pi^m = |\phi^m\rangle\langle\phi^m|$. We show that G supports satisfying product states on open neighborhoods in the projector manifold if it has a qubit-clause dimer covering that covers all of the clauses – *ie.* an assignment of a unique qubit to each clause (see Fig. 1). Appealing to the algebraic properties of the space of projectors, this shows that G is in fact PRODSAT for all choices of Π^m . It also suggests a counting criterion for the ground state degeneracy. Thus, we will show

Theorem 1. *If an interaction graph G , viewed as a bipartite factor graph of qubits and clauses, has a dimer covering of its clauses, then*

1. *there exist open neighborhoods $U \subset (\mathbb{CP}^{2^k-1})^M$ of the rank 1 projector manifold such that $\Pi^m \in U \Rightarrow (G, \Pi^m)$ is PRODSAT.*
2. *Moreover, (G, Π^m) is PRODSAT for all $\Pi^m \in (\mathbb{CP}^{2^k-1})^M$.*

As a practical matter, product satisfiability is a question about the existence of solutions to a collection of M homogeneous equations in $N \mathbb{CP}^1$ degrees of freedom (unless otherwise noted, paired i indices contract, but m indices do not):

$$\phi_{i_1 i_2 \dots i_k}^m \psi_{m_1}^{i_1} \psi_{m_2}^{i_2} \dots \psi_{m_k}^{i_k} = 0 \quad \forall m \in G \quad (2)$$

where $\psi_{m_1}^{i_1}$ is the i_1 'th component of the state on the m_1 'th qubit and ϕ^{m*} is the state onto which Π^m projects. Since we are interested in generic choices of the ϕ^m , one might expect naive constraint counting to determine whether $(M \leq N)$ or not $(M > N)$ a given graph G is generically PRODSAT. Unfortunately, things are not quite so simple because a G with a low average density $(M \leq N)$ may still have subgraphs that are dense enough to overconstrain the equations.

To make this last observation sharper, let us briefly restrict our attention to product projectors. In this case, Eq. (2) reduces to

$$(\phi_{i_1}^{m_1} \psi_{m_1}^{i_1})(\phi_{i_2}^{m_2} \psi_{m_2}^{i_2}) \dots (\phi_{i_k}^{m_k} \psi_{m_k}^{i_k}) = 0 \quad \forall m \in G \quad (3)$$

This equation is clearly satisfied if any of its factors is satisfied, which requires at least one of the k qubits to be (uniquely) fixed orthogonal to the local projector. Thus, for generic ϕ of product form, in which the local projectors on a site n are not parallel, the full set of M equations is solvable if and only if one can uniquely associate a qubit n to each clause m . Graphically, this is equivalent to the existence of a dimer covering of the bipartite factor graph of G which covers all M clauses. Note that we refer to matchings as dimer coverings if each clause is uniquely associated to a qubit even if some qubits are left unpaired. See Fig 1.

We now relax the constraint that ϕ take product form and argue that the existence of the above dimer covering is still the relevant constraint counting for product satisfiability. Let us consider an instance (G, ϕ) and assume that it is PRODSAT; *w.l.o.g.* we may choose local bases such that $|\Psi\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \otimes \dots \otimes \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ is a satisfying assignment. Choosing stereographic coordinates $\begin{pmatrix} z_n \\ 1 \end{pmatrix}$ for each qubit, we find the constraint equation is:

$$\phi_{i_1 i_2 \dots i_k}^m z_{m_1}^{i_1} z_{m_2}^{i_2} \dots z_{m_k}^{i_k} = 0 \quad (4)$$

where now superscript i_1 indicates exponentiation and the i indices run from 0 to 1. Since $z_n = 0$ satisfies this equation, we find that $\phi_{00\dots 0}^m = 0$ necessarily but that the remaining components of the ϕ^m are unconstrained.

Now we perturb $\phi^m \mapsto \phi^m + \delta\phi^m$ and attempt to follow the solution $z_n = 0 \mapsto z_n + \delta z_n$. To linear order, we have (for each clause $m \in G$):

$$\phi_{10\dots 0}^m \delta z_{m_1} + \phi_{01\dots 0}^m \delta z_{m_2} + \dots + \phi_{00\dots 1}^m \delta z_{m_k} = -\delta\phi_{00\dots 0}^m \quad (5)$$

This is a collection of M sparse linear inhomogeneous equations for the N variables δz_n . Let Ω_{mn} be the $M \times N$ matrix of coefficients on the left hand side. This matrix connects qubits to clauses and is sparse in the same pattern as the node-edge adjacency matrix A_{mn} of G .

Equation (5) is solvable for arbitrary $\delta\phi_{00\dots 0}^m$ if and only if the matrix Ω is surjective. Trivially, this requires $M \leq N$, but it also requires that some $M \times M$ subdeterminant of Ω be nonzero. Let us define a polynomial discriminant of surjectivity as the tuple of all possible subdeterminants of Ω :

$$\text{disc}_\tau(\Omega) = \sum_{\sigma \in S_M} (-1)^\sigma \Omega_{1\tau(\sigma(1))} \dots \Omega_{M\tau(\sigma(M))} \quad (6)$$

where $\tau : M \hookrightarrow N$ runs over all injective maps. For formal components ϕ , $\text{disc}_\tau(\Omega)$ is not identically zero if and only if one of the terms in the above sums is not identically zero; each non-zero term corresponds to a particular pairing of clauses and qubits in the adjacency graph which covers all M clauses but need not cover all of the

nodes. Thus if G has such a dimer covering, Ω is almost always surjective, whereas it is never surjective otherwise.

We can summarize the above analysis as follows: for G having dimer coverings, generic choices of ϕ for which (G, ϕ) is PRODSAT may be extended to open neighborhoods around ϕ which are PRODSAT. Moreover, G has product states at generic (non-parallel) *product* projectors. One can repeat the perturbative analysis of Eq. (5) near such product projectors and show that the dimer covering condition allows the same extension onto open neighborhoods in the full projector space. This proves part 1 of theorem 1. [16]

The above analysis is essentially local, showing that the set $W \subset (\mathbb{CP}^{2^k-1})^M$ of PRODSAT projector choices is nonempty and full dimension; to extend product satisfiability to the *entire* projector manifold, we need to appeal to some basic results in complex projective geometry: 1. Any Zariski-closed full dimension subspace of an irreducible complex projective space (such as $(\mathbb{CP}^{2^k-1})^M$) is in fact the whole space and 2. W is Zariski-closed. This shows part 2 of theorem 1. For more details, see Appendix A.

Finally, we note that the converse statement also holds (also shown in Appendix A):

Theorem 2. *If G does not have dimer coverings, (G, Π^m) is not PRODSAT for almost all Π^m .*

A. Counting product states at $M = N$

We now comment on the marginal case when $M = N$. In this case, one expects the satisfying product states, when they exist, to be discrete. Let us return to projectors of product form and consider a choice ϕ_0 in which no local projectors on a single site n are parallel, as mentioned above. In this case, it is clear that there is a one-to-one correspondence between dimer coverings c of G and product state solutions ψ_c . The number of product states at the special point ϕ_0 is given by the number of dimer coverings.

Now, one can pick any of the ψ_c and repeat the local perturbative analysis around ϕ_0 . This will show a unique way to continuously extend ψ_c in an open neighborhood of ϕ_0 . Thus, at least on open neighborhoods within projector space, the number of product states is bounded below by the number of dimer coverings. Since the problematic points in ϕ space for this local analysis are actually degenerate points with higher satisfiability, this should be a lower bound on the number of product states for any ϕ and it should be the exact counting almost everywhere.

With a handle on all of these product states, one might hope to make some sharp statements about the QSAT dimension of G – that is, the generic $R_G = |\ker(H)|$ for the graph G . Let $PS = \text{Span}\{\psi_c\}$ be the span of the dimer covering states at some generic product projector point ϕ_0 . If $R_{PS} = |PS|$ of the ψ_c are linearly independent,

they will remain so on an open neighborhood of ϕ_0 , since the determinant of their overlap matrix is a smooth function. We recall that the QSAT dimension R_G takes its minimal value *almost everywhere* on projector space [5]. Since R_{PS} lower bounds R_G on an open neighborhood, it must therefore lower bound R_G for all projectors ϕ .

Using the dimer covering states to characterize the full SAT subspace $\ker(H)$ for a given graph G requires answering two questions: 1. Are the product states linearly independent? and 2. Do they span $\ker(H)$? For some simple families of graphs, it is possible to prove the linear independence of the dimer states by exploiting the geometry of their dimer coverings. On the other hand, it is easy to construct graphs where the dimer states are not independent: for a fully connected bipartite factor graph, the number of dimer coverings is $N!$, which is significantly greater than the size 2^N of the Hilbert space.

The second question is harder to address analytically. Product states do span the kernel for many graphs that we have studied numerically, but we also know examples of small graphs for which they do not. Rather, R_G can be strictly greater than the number of dimer coverings. Indeed, H restricted to product projectors may have an even greater satisfying dimension than it does for general entangled projectors, even though the number of product states should be no greater.

III. APPLICATION TO RANDOM GRAPHS

From the previous section, we know that the existence of dimer coverings on G implies quantum satisfiability – and that their nonexistence shows satisfying product states are no longer generic. For the random graph ensemble, we can therefore lower bound the QSAT transition α_c by the threshold for the presence of dimer coverings α_{dc} . This turns out to be precisely the point at which the hypercore of the random graph reaches the critical density $M_{hc} = N_{hc}$, beyond which it is certainly impossible to dimer cover the clauses. In general, the hypercore emerges at $\alpha_{hc} < 1$ with a clause density less than 1 and as α grows, it gets denser. Only at $\alpha_{dc} > \alpha_{hc}$ does it become critically constrained. See Table I for numeric values of α_{dc} .

In order to show that dimer coverings exist below α_{dc} , we rely on a somewhat indirect argument based on the results of Mezard *et al.* [9] regarding the random k -XORSAT problem, or, in physics parlance, the p -spin Ising glass (with $p = k$). In this problem, N classical bits live on an interaction graph G whose (hyper-)edges define parity check constraints. That is, for each edge m , we must choose $J_m \in \{0, 1\}$ and then determine whether the following linear constraint equations are satisfiable:

$$x_{m_1} + x_{m_2} + \cdots + x_{m_k} = J_m \quad (7)$$

where x_n are the bits and all arithmetic takes place in the field \mathbb{Z}_2 .




FIG. 2: Phase diagram of random k -QSAT for $k = 3$. The labeled transitions are (left to right): the emergence of the giant component (α_{gc}); the emergence of the hypercore (α_{hc}); the disappearance of product states ($\alpha_{ps} = \alpha_{dc}$); the best current upper bound on satisfiability due to Bravyi *et al*[7] (α_c^+); the classical cavity transition for k -SAT (α_{cav}) and the weak Pauling bound (α_{wb}). At large k , the known upper bounds all scale exponentially while the best lower bound (α_{ps}) is bounded above by $\alpha = 1$.

This defines a linear system of inhomogeneous equations of exactly the same form as what we discovered when linearizing PRODSAT:

$$A_{mn}x_n = J_m \quad (8)$$

where A is the node-edge adjacency matrix of G . In this case, A is surjective if and only if all choices of J_m result in satisfiable instances of XORSAT. Moreover, if A is surjective over \mathbb{Z}_2 , so is the sparse matrix $\Omega_{mn} = \phi_{mn}A_{mn}$ over \mathbb{C} with arbitrary components ϕ_{mn} : if Ω were not generically surjective, the rank discriminating polynomials $\text{disc}_\tau(\Omega)$ of Eq. 6 would be identically zero, including when mapped homomorphically to \mathbb{Z}_2 .

Mezard *et al* show that the random XORSAT problem is unfrustrated for $\alpha < \alpha_{dc}$. Indeed, they show precisely that with probability going to 1 in the thermodynamic limit, the random graph G is XORSAT for *all* choices of J_m . In particular, this implies that A is surjective for $\alpha < \alpha_{dc}$ and therefore that Ω is as well, dimer coverings exist and the QSAT problem is satisfiable. Meanwhile, for $\alpha > \alpha_{dc}$, $M > N$ on the hypercore and the dimer coverings are gone.

We note that the previously known lower bound was given by α_{hc} , the threshold for the emergence of a hypercore in G . As k becomes large, α_{hc} approaches zero while α_{dc} approaches 1 (from below). This significantly raises the known lower bound for quantum satisfiability, but it still does not approach the best known upper bounds, which scale with 2^k . See the phase diagram in Fig. 2 for the numerical values at $k = 3$.

IV. BEYOND PRODSAT

In the previous section, we argued that random k -QSAT is PRODSAT for all k , at clause densities $\alpha < \alpha_{dc}(k) = \alpha_{ps}(k)$ with $\alpha_{ps}(2) = \frac{1}{2}$, $\alpha_{ps}(3) = .92\dots$ and $\alpha_{ps}(k \rightarrow \infty) \rightarrow 1^-$. (A heuristic way of understanding the latter limit is provided by the observation that for large k , the probability of having sites with low coordination is suppressed exponentially $\sim e^{-\alpha k}$, so that the density of the random graph and its 2-core are almost the same.)

k	Quantum			Classical	
	α_{hc}	α_{ps}	α_c^+	α_{cl}^-	α_{cl}^+
2	0.5	0.5	0.5	1	1
3	0.82	0.92	3.59	3.52	4.49
4	0.77	0.98	7.98	7.91	10.23
5	0.70	0.99	16.00	18.79	21.33
∞	0^+	1^-	$2^{k-1} \ln 2$	$2^k \ln 2$	$2^k \ln 2$

TABLE I: Summary of critical values for various random k -(Q)SAT properties. The values for the emergence of a hypercore α_{hc} and for the dimer coverability/product satisfiability $\alpha_{ps} = \alpha_{dc}$ of the random interaction graph may be found in [9]. $\alpha_c^+(3)$ comes from the BMR ‘nosegay’ bound[7] while the values at larger k and in the large k limit come from the ‘sunflower’ bound (cf. Appendix C). For comparison, we include the best known rigorous bounds on the classical satisfiability transition [10, 11].

What happens for $\alpha > \alpha_{ps}(k)$? Reliable information on this only exists for $k = 2$: as the ensemble goes unPRODSAT, it immediately goes unSAT as well. For $k \geq 3$, no such result has as yet been established. Rather, the best upper bounds on the extent of the SAT phase, due to BMR[7], are well separated from the PRODSAT transition: a detailed analysis of the equations they derive yields the upper bounds of Table I. The large- k asymptote of this bound is $\alpha_c^+(k) = 2^{k-1} \ln 2 + O(\frac{1}{\ln k})$, which is *exponential* in k , leaving plenty of space above $\alpha_{ps} = 1^-$ for a SAT–unPRODSAT phase.

In the following, we adduce some numerical pointers regarding the existence of such a phase, which is arguably among the most interesting open questions in this problem.

A. Direct search for SAT–unPRODSAT instances

First, we have directly searched for instances – irrespective of their probability in the random graph ensemble – which are SAT–unPRODSAT. Our search for satisfying assignments proceeded via (i) exact diagonalisation of systems with up to 14 qubits and (ii) a numerical nonlinear equation solver for finding product solutions. Our results were all consistent with the picture of product states outlined above. In addition, we found graphs with satisfying assignments with $M > N$. In these cases method (ii) failed to turn up any product ground states. We also diagnosed directly the presence of non-product ground states by considering the rank of the reduced density matrix corresponding to a ground state after an arbitrary portion of the system is traced over, cf. Fig. 3. The resulting rank of the density matrix is bounded above by the number of zero-energy states, provided these are spanned by product states (c.f. Appendix B for a proof). As an illustration, we show the corresponding rank for the 2-SAT problem in Fig. 3(a) for reduced density matrices of up to five qubits from instances containing 13 qubits. While the SAT instances show the expected rank

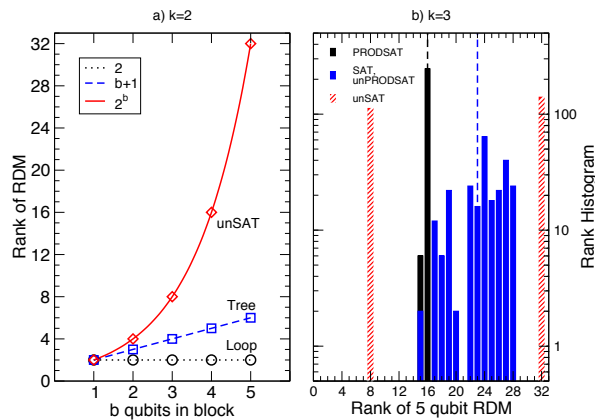


FIG. 3: (Color online) Rank analysis of reduced density matrices (RDM) in a ground state: (a) $k=2$: Rank of the reduced density matrix containing $b = 1, \dots, 5$ qubits for different types of interaction graphs. (b) $k=3$: Histogram of ranks of $b = 5$ reduced density matrices for all 252 different partitions ($N = M = 10$) for three different instances.

for product states, the unSAT instance has full rank, suggesting generically entangled wavefunctions in the $k=2$ unSAT phase. In Fig. 3(b) we performed a similar analysis for $k=3$. We determined rank histograms for all 252 five-qubit reduced density matrices of $N = M = 10$ systems for the three different instances listed in [17]. The PRODSAT instance leads to a histogram where the rank is bounded by the number of satisfying ground states (16 for this instance). The SAT-unPRODSAT instance shows ranks which are both smaller and larger than the number of satisfying ground states (23 for this instance), but do not reach the maximum of $2^5 = 32$. Finally, the unSAT instance has full rank 32 on the two-core (the rank 8 occurring for some partitions is due to a satisfied dangling clause). We do indeed find that, in the presence of $M > N$ on a (sub)graph, entangled states necessarily appear in any basis of the ground states.

B. Finite size scaling

Where is the the SAT/unSAT transition α_c ? Numerical studies on finite-site systems offer limited information as the properties of small random graphs are quite different from those of large ones. We do consistently find the presence of SAT instances with $N > M$: for 3-SAT, we investigated 101 random graphs for sizes with up to 13 sites using complete diagonalization and up to 20 sites using an iterative diagonalization technique, see left panel of Fig. 4. While complete diagonalization is limited by the growing matrix size but otherwise numerically stable, the iterative Lanczos procedure faces challenges when resolving small energy gaps occurring at the SAT-unSAT transition. This leads to some instances becoming undecidable, i.e. not converging within several thousand iterations. The presence of undecidable instances is denoted




FIG. 4: (Color online) Probability to find SAT instances among 101 random graphs for (a) $k=3$ and (b) $k=4$. Numerical results obtained by complete diagonalization for $N \leq 13$ and iterative diagonalization for $N = 14, 16$ ($k=3, 4$) and 18, 20 ($k=4$). For the latter system sizes filled symbols indicate the occurrence of (numerically) undecidable instances.




FIG. 5: (Color online) Finite size scaling of $\alpha_c(N)$ for $k=3, 4$ obtained from the data displayed in Fig. 4. The empty symbols denote α values where the SAT ratio is 0.5. The lower (upper) ticks correspond to the largest α values with SAT ratio > 0.9 (smallest α values with SAT ratio < 0.1).

by filled symbols [18]. For 3-SAT, we find that curves for larger N nicely converge, suggesting $\alpha_c \approx 1 \pm 0.06$, where the error bar is estimated from the largest N in Fig. 5. For 4-SAT, instances are SAT for larger values of $\Delta = M - N$ (for fixed N), even though Δ appears to decrease with N (inset in right panel of Fig. 4). The finite size scaling shown in Fig. 5 is also compatible with $\alpha_c \approx 1$, but with a larger uncertainty of ≈ 0.2 .

Thus, we find many examples of generically SAT-unPRODSAT graphs G and even some evidence that the unSAT phase transition in the random ensemble lies at $\alpha > \alpha_{ps}$.

V. IN CLOSING

In this paper we have reported considerable progress in understanding product satisfiability. Applied to the random graph ensembles introduced in [5], it yields a defini-

tive lower limit to the clause density at which satisfiability necessarily requires entangled states and certainly an improved lower bound to the location of the SAT-unSAT phase transition. Between the work reported in this paper and by BMR [7], there has thus been a fair amount of progress in narrowing the location of this phase transition from the initial work [5].

Nonetheless, there remains a large gap between the best lower bound $\alpha_c^- = \alpha_{ps} = O(1)$ and the best upper bound $\alpha_c^+ = O(2^k)$. By contrast, the classical random k -SAT problem has tight bounds on its satisfiability transition that scale as 2^k [10]. It is not obvious how to translate the techniques used in proving the tight lower bounds for the classical problem to the quantum case.

Ambainis and co-workers[8] had recently proven a quantum version of the Lovász local lemma, which directly shows the quantum satisfiability of graphs G of bounded degree. Due to this degree restriction, the lemma does not directly apply to the Erdős random graph ensemble. Thus, they have used the characterization of product satisfiability shown here coupled with the quantum Lovász lemma in order to show that a wide SAT-unPRODSAT phase in the parameter α exists: the SAT regime extends to $\alpha_c > 2^k/(12 \cdot e \cdot k^2)$ which, for $k \geq 13$, is significantly beyond the PRODSAT transition at $\alpha_{ps} \approx 1$.

Many other questions remain open and it is to these that we turn in closing:

1. Does the SAT-unPRODSAT phase exist for all $k \geq 3$ or does it really only arise at some larger k ?
2. Does an easily checked classical graph property encode generic quantum satisfiability? Does a straightforward function compute the generic dimension R_G ?
3. What can be said about higher rank projectors? The product satisfiability condition translates straightforwardly to the higher rank case by simply treating a rank r projector as a stack of rank 1 projectors. Thus, higher rank QSAT instances may naturally yield more entangled ground states than the rank 1 case. Indeed, preliminary work indicates that this is the case[12].

Acknowledgements

We thank S. Bravyi, C. Moore, T. Osborne, D. Nagaj, A. Obus and especially O. Sattath for many enlightening conversations. We would also like to thank the generous hospitality of the Erwin Schrödinger Institute in Vienna where parts of this work were completed.

Appendix A: Formal PRODSAT proof

We expand on the complex algebraic details of theorems 1 and 2. That is, we show that interaction graphs G with dimer coverings are PRODSAT *for any* choice of projectors $\{\Pi^m\}_{m=1}^M$ and conversely that G without dimer coverings are not PRODSAT *for almost all* such choices. The relevant algebraic geometry background can be found in [13, 14].

Proof. Let $P = (\mathbb{CP}^{2^k-1})^M$ be the space of k -qubit projectors, $S = (\mathbb{CP}^1)^N$ be the space of product states, and $V \subset P \times S$ be the space of pairs of $(\{\Pi^m\}_{m=1}^M, \{|\psi_n\rangle\}_{n=1}^N)$ such that $|\Psi\rangle = |\psi_1\rangle \otimes \cdots \otimes |\psi_N\rangle$ satisfies the projectors Π^m . We note that P , S and $P \times S$ are each irreducible complex projective spaces in the Zariski topology and that V , as the zero-locus of collection of homogeneous algebraic constraints, is a closed subspace of $P \times S$. Let $\pi : P \times S \rightarrow P$ be the projection map onto the first component. Since S is projective, this map is closed.

We are interested in the set of projectors $W \subseteq P$ which are satisfiable by product states. These are precisely the projectors in the image of V under the projection π :

$$W = \pi(V) \quad (\text{A1})$$

Since π is closed, W is closed. By the local arguments of Section II, we know that W is full dimension and nonempty. Since P is irreducible and projective, it has no proper full dimension closed subspaces; *ie.* $W = P$.

Conversely, if G does not have dimer coverings, there exist choices of product projectors Π^m which are clearly unsatisfiable by product states (see Eq. (3)). Hence W is a closed proper subspace of P and cannot be full dimension. Thus, W has codimension at least 1 in P and almost all choices of projector Π^m will fail to be product satisfiability. \square

Appendix B: Reduced density matrix diagnostic for product states

Let us consider a graph Λ and a wavefunction $|\Psi\rangle$, which is expressible as a sum of wavefunctions each of which is a product wavefunction on the sites $i \in \Lambda$. Let $\gamma \subset \Lambda$, and let $\rho = |\Psi\rangle\langle\Psi|$ and $\tilde{\rho}_\gamma = \text{Tr}_{\Lambda \setminus \gamma} \rho$. Then $\text{rank } \tilde{\rho}_\gamma \leq N_0^{(\gamma)} \leq N_0$ where N_0 is the number of linearly independent product states, and $N_0^{(\gamma)}$ the dimension of their span when restricted to γ .

Let us label a product basis for the ground state manifold by: $\{\Phi_\alpha | \alpha = 1 \dots N_0\}$ with $|\Phi_\alpha\rangle = \otimes_{i=1}^{n_s} |\phi_i^\alpha\rangle$, $\langle\phi_i^\alpha | \phi_i^\alpha\rangle = 1$. Here $|\Lambda| = n_s$. Thus, $\langle\Phi_\alpha | \Phi_\alpha\rangle = 1$ (normalised) and $\langle\Phi_\alpha | \Phi_\beta\rangle = \prod_{i=1}^{n_s} \langle\phi_i^\alpha | \phi_i^\beta\rangle$ (not necessarily orthogonal).

A general ground state is

$$|\Psi\rangle = \frac{1}{N} \sum_{\alpha=1}^{N_0} \lambda_\alpha |\Phi_\alpha\rangle \quad (\text{B1})$$

with

$$\aleph^2 = \sum_{\alpha, \beta=1}^{N_0} \lambda_\alpha \lambda_\beta^* \prod_{i=1}^{n_s} \langle \phi_\beta^i | \phi_\alpha^i \rangle, \quad (\text{B2})$$

for which

$$\rho = |\Psi\rangle\langle\Psi| = \frac{1}{\aleph^2} \sum_{\alpha, \beta=1}^{N_0} \lambda_\beta^* \lambda_\alpha |\Phi_\alpha\rangle\langle\Phi_\beta|. \quad (\text{B3})$$

Tracing over the qubits on subset $\Lambda \setminus \gamma$ of the set Λ yields:

$$\tilde{\rho}_\gamma = \frac{1}{\aleph^2} \sum_{\alpha, \beta=1}^{N_0} \lambda_\beta^* \lambda_\alpha [\otimes_{i' \in \gamma} |\phi_{i'}^\alpha\rangle] \left[\otimes_{j' \in \gamma} \langle \phi_{j'}^\beta | \right] \prod_{k' \in \Lambda \setminus \gamma} \langle \phi_{k'}^\beta | \phi_{k'}^\alpha \rangle. \quad (\text{B4})$$

$\tilde{\rho}_\gamma$ is at most an $N_0 \times N_0$ matrix, with at most rank N_0 irrespective of the choice of γ .

If $\dim(\text{span}\{\otimes_{i' \in \gamma} |\phi_{i'}^\alpha\rangle | \alpha = 1 \dots N_0\}) = N_0^{(\gamma)} < N_0$, then we can write, for $\alpha = N_0$ (say): $\otimes_{i' \in \gamma} |\phi_{i'}^{N_0}\rangle = \sum_{\alpha=1}^{N_0-1} \mu_\alpha^{(N_0)} [\otimes_{i' \in \gamma} |\phi_{i'}^\alpha\rangle]$, so that terms like

$$\left[\otimes_{i' \in \gamma} |\phi_{i'}^{N_0}\rangle \right] \left[\otimes_{j' \in \gamma} \langle \phi_{j'}^{N_0} | \right] = \sum_{\alpha, \beta=1}^{N_0-1} \mu_\alpha^{(N_0)} \mu_\beta^{*(N_0)} [\otimes_{i' \in \gamma} |\phi_{i'}^\alpha\rangle] \left[\otimes_{j' \in \gamma} \langle \phi_{j'}^\beta | \right], \quad (\text{B5})$$

take on the same form as the ones in (B4). Using such substitutions repeatedly, one finds that the reduced dimensionality manifestly has a rank of at most $N_0^{(\gamma)}$.

In particular, for the case of 2-QSAT and Λ being a tree ($||\Lambda|| + 1 = N_0$), $N_0^{(\gamma)} = ||\gamma|| + 1$. When a qubit being traced out is a dangling one, this result is obvious. In the case of an internal qubit (labelled 2, say), we note that the Bravyi construction on Λ induces one on γ via $\top_{13}^{(\gamma)} = \top_{12} \top_{23}$, where \top_{ij} represents the Bravyi transfer matrix between for the projector ij , so that the resulting state counting is that of the smaller graph.

Appendix C: Improved estimate of α_c^+ from BMR sunflowers

In this section we provide an integral representation of the upper bound on the ground state dimension obtained in [7] by studying ‘sunflower’ graphs for generic k . By means of this we can obtain upper bounds $\alpha_c^+(k)$ for $k \geq 4$ and an asymptotic estimate for large k . From [7] we can read an upper bound to the dimension of the satisfying space R_G as

$$\lim_{n \rightarrow \infty} \frac{1}{n} R_G \leq \ln 2 + \sum_{d=0}^{\infty} a_d \left(d \ln(1 - 2^{-k+1}) + \ln \left(1 + \frac{d}{2^k - 2} \right) \right) \equiv S, \quad (\text{C1})$$

where

$$a_d = \int_0^1 dt \frac{(k\alpha t^{k-1})^d}{d!} e^{-k\alpha t^{k-1}} \equiv \int_0^1 dt a_d(t). \quad (\text{C2})$$

By inverting the sum over d and the integral over t we find

$$S = \ln 2 + \alpha \ln(1 - 2^{-k+1}) + \int_0^1 dt \langle \ln \left(1 + \frac{d}{2^k - 2} \right) \rangle_t. \quad (\text{C3})$$

The last is a t -dependent average with the Poisson distribution $a_d(t)$. We use the representation

$$\ln(1+x) = \int_0^\infty ds \frac{e^{-s}}{s} (1 - e^{-sx}) \quad (\text{C4})$$

to perform the average over d first (using the generating function of the Poisson distribution) and the integral over t afterwards. We are left with

$$S = \ln 2 + \alpha \ln(1 - 2^{-k+1}) + \int_0^\infty ds \frac{e^{-s}}{s} \left(1 - \frac{q^{1/(k-1)}}{k-1} \left(\Gamma\left(\frac{1}{k-1}\right) - \Gamma\left(\frac{1}{k-1}, q\right) \right) \right) \quad (\text{C5})$$

where $q = k\alpha(1 - e^{-s/(2^k-2)})$ and $\Gamma(z, x)$ is the incomplete gamma function (with $\Gamma(z, 0) = \Gamma(z)$). This integral representation, although seemingly complicated, is perfectly suited for numerical integration, since the integrand is always finite and cut-off at $s = O(1)$, beyond which it decreases exponentially in s .

One can obtain the numbers quoted in the text for the threshold $\alpha_c^+(k)$ to arbitrary accuracy using Mathematica. Moreover in the limit $k, \alpha \rightarrow \infty$ and $\alpha/2^k$ finite, relevant for the large k asymptotics, we can bound the value of the integral, solve for $S = 0$ and find the quoted result

$$\alpha_c^+ = 2^k \frac{\ln 2}{2} + O(1/\ln k). \quad (\text{C6})$$

it is worth noting that this bound is half of the value of the *lower* bound of k -SAT which is known to become

exact in the large k limit.

We have not been able to analyze in the same detail the upper bound given by the ‘nosegay’ graphs of [7]. Although this bound would be stricter for *any* k than that coming from sunflower graphs, our preliminary results point towards the same asymptotic scaling (C6). It would be interesting to understand whether for large k this bound is tight, as happens in classical random k -SAT.

-
- [1] M. R. Garey and D. S. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Series of Books in the Mathematical Sciences (W. H. Freeman, New York, 1979).
 - [2] S. Bravyi, e-print arXiv:quant-ph/0602108.
 - [3] S. Kirkpatrick and B. Selman, *Science* **264**, 1297 (1994).
 - [4] R. Monasson, R. Zecchina, S. Kirkpatrick, B. Selman, and L. Troyansky, *Nature* **400**, 133 (1999).
 - [5] C. R. Laumann, R. Moessner, A. Scardicchio, and S. L. Sondhi, *Quantum Inf. Comput.* **10**, 0001 (2010).
 - [6] M. Molloy, *Random Struct. Algorithms* **27**, 124 (2005).
 - [7] S. Bravyi, C. Moore, and A. Russell, e-print arXiv:0907.1297v2.
 - [8] A. Ambainis, J. Kempe, and O. Sattath, in *42nd Annual ACM Symposium on Theory of Computing* (2009), (in press), e-print arXiv:0911.1696v1.
 - [9] M. Mézard, F. Ricci-Tersenghi, and R. Zecchina, *J. Stat. Phys.* **111**, 505 (2003).
 - [10] D. Achlioptas, *Random Satisfiability* (IOS Press, 2009), vol. 185 of *Frontiers in Artificial Intelligence and Applications*, chap. 8, pp. 245–270, ISBN 978-1-58603-929-5.
 - [11] J. Diaz, L. Kirousis, D. Mitsche, and X. Perez-Gimenez, in *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2008)*, edited by R. Hariharan, M. Mukund, and V. Vinay (Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany, Dagstuhl, Germany, 2008), vol. 2 of *Leibniz International Proceedings in Informatics*.
 - [12] R. Movassagh, E. Farhi, J. Goldstone, D. Nagaj, T. J. Osborne, and P. W. Shor, e-print arXiv:1001.1006.
 - [13] J. S. Milne, *Algebraic geometry (v5.20)* (2009), www.jmilne.org/math/.
 - [14] D. Mumford, *Algebraic Geometry I: Complex Projective Varieties*, Classics in Mathematics, Vol. 1 (Springer, New York, 1995).
 - [15] This is the class of decision problems which can be verified efficiently on a quantum computer with perfect completeness – *ie.* in which the verifier never rejects a valid proof. QMA₁ is closely related to the somewhat more general class QMA, which does not require perfect completeness.
 - [16] For those who prefer a more rigorous looking style, we note that the first order perturbative calculation performed in this section is identical to checking the conditions under which the implicit function theorem provides a smooth map from projector space to the manifold of satisfying product states near a known solution.
 - [17] The three $N=M=10$ graph instances are listed here for convenience: a) PRODSAT: $\{\{3,5,8\}, \{5,6,7\}, \{0,6,8\}, \{1,3,5\}, \{0,2,5\}, \{1,3,9\}, \{1,4,9\}, \{0,1,5\}, \{2,3,7\}, \{0,1,2\}\}$ b) SAT-unPRODSAT: $\{\{0,1,5\}, \{0,4,5\}, \{2,5,8\}, \{5,7,9\}, \{5,6,7\}, \{5,7,8\}, \{6,8,9\}, \{0,3,5\}, \{4,6,8\}, \{2,4,9\}\}$ c) unSAT: $\{\{2,7,8\}, \{0,4,7\}, \{4,5,6\}, \{1,5,6\}, \{1,6,9\}, \{3,5,7\}, \{0,3,7\}, \{5,6,7\}, \{1,3,7\}, \{1,6,7\}\}$
 - [18] Among the 101 graphs we found the following maximum number of undecidable instances: $N = 14, k = 3$, max 4 undecidable; $N = 16, k = 3$, max 9 undecidable; $N = 14, k = 4$, max 24 undecidable; $N = 16, k = 4$, max 57 undecidable; $N = 18, k = 4$, max 62 undecidable; $N = 20, k = 4$, max 58 undecidable.